## REMARKS

Claims 1-50 are pending in the present application. Claims 1, 2, 24, 25 and 48 have been amended herewith. Reconsideration of the claims is respectfully requested.

### I. 35 U.S.C. § 112, 2nd Paragraph

The Examiner rejected Claims 2, 25 and 48 under 35 U.S.C. § 112, 2nd paragraph as being vague and indefinite. Applicants have amended Claims 1, 2, 24, 25 and 48 in response thereto, to clarify that the claimed system resource shares and tier are characteristics of each class (Claims 1 and 24), and that a percentage of system resources to be allocated to the process is based upon the system resource shares of the class to which the process was classified as well as the system resource shares associated with the other classes (Claims 2, 25 and 48).

Therefore, the rejection of Claims 2, 25 and 48 under 35 U.S.C. § 112, 2nd paragraph has been overcome.

### II. 35 U.S.C. § 103, Obviousness

The Examiner rejected Claims 1- 50 under 35 U.S.C. § 103 as being unpatentable over Ferguson et al. (U.S. Patent 5,504,894) in view of Bhatti et al. (U.S. Patent 6,304,906). This rejection is respectfully traversed.

To establish prima facie obviousness of a claimed invention, all of the claim limitations must be taught or suggested by the prior art. MPEP 2143.03. *See also, In re Royka*, 490 F.2d 580 (C.C.P.A. 1974) (emphasis added by Applicants). In the absence of a proper prima facie case of obviousness, an applicant who complies with the other statutory requirements is entitled to a patent. *See In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992). Applicants will now show that all of the claim limitations are not taught or suggested by the cited references, and thus Applicants are entitled to a patent for the claims recited in the present application.

With respect to Claim 1 (and dependent claims thereof), Applicants show that none of the cited references teach or suggest a class having *both* system resource shares and a tier. This is shown in the preferred embodiment in Applicants' Figure 6 at 670 and 680. By providing classes having both system resource shares and a tier, it is possible to

associate a given process with a particular class such that the attributes of a process may be used to perform a classification of the process into a defined class having *both* system resource shares and a tier. In addition, by providing classes having *both* system resource shares and a tier, it is possible to share system resources amongst processes classified to a given class (Specification page 12, lines 10-30), as well as favor processes in a given class tier over processes in another class tier (Specification page 13, lines 3-19). This two-fold class characterization thus allows processes to be more effectively classified in a finer granularity than what was previously possible in the prior art.

The cited Ferguson reference teaches receipt of a transaction that belongs to a pre-defined transaction class (Col. 4, lines 26-27). The only characteristic associated with the Ferguson transaction classes is average response time goals (Col. 5, lines 24-26). Thus, the cited Ferguson reference merely teaches a class having a single characteristic, and this single characteristic is not (i) system resource shares, or (ii) a tier, as claimed.

The cited Bhatti reference teaches prioritization of access requests into classes such that preferential treatment can be given to some of the users accessing the data service system (Col. 1, lines 62-65). A user access request classification system includes a request classifier that classifies each of the access requests into one of a plurality of classes based on a predetermined classification policy (Col. 2, lines 4-7). The user access request classification system also includes a plurality of buffers coupled to the request classifier. Each buffer corresponds to one of the classes to receive the access requests classified by the request classifier as belonging to that class such that the data service system processes the access requests in accordance with their classes (Col. 2, lines 7-13). Use of the buffers negates any requirement or desire to having two-fold class characteristics as a received request is placed in an appropriate buffer/queue for subsequent processing. In effect, the received requests are segregated into hardware buffers and therefore there is no need for classes having two-fold characteristics of both system resource shares and a tier, as claimed.

While Bhatti does mention a two-tier or class scheduling policy, this policy specifies *which of the queues* should be first serviced and *which request* in the queue should be first serviced (Col. 8, lines 57-62). This occurs after a *request* has already been placed in a queue, and thus has nothing whatsoever to do with classifying a *process* into a

given class having system resource shares and a tier. Bhatti merely teaches processing of a given request, and then the processing of another given request. Applicants have amended Claim 1 to further clarify this distinction between requests (as taught by Bhatti) and concurrently executing processes (as claimed).

The only mention by Bhatti of process classification is a single-fold priority classification. As described by Bhatti at Col. 7, lines 12-24:

> "In addition, the server application 53 may also include a modifier 54. The modifier 54 is used to *mark the process or thread that executes the request* received by the server application 53 from one of the access request classification systems 52-52n. *The modifier 54 marks the process or thread with a priority that corresponds to the importance of the request to be processed.* The modifier 54 can be implemented by known technology. For example, the modifier 54 can be implemented using the HP PRM (Process Resource Management) technology available from Hewlett-Packard Co. of Palo Alto, Calif. Another example of implementing the modifier 54 is using the known Unix based NICE technology." (emphasis added by Applicants)

Thus, Bhatti merely teaches a *single-fold* priority characteristic for a process. In contrast, the claimed invention is directed to classifying processes into a class, where the class has *two-fold* characteristics (system resource shares and tier). As every element of Claim 1 is not taught or suggested by the cited reference, it is shown that Claim 1 (and dependent claims thereof) is not obvious in view of the cited references.

Applicants traverse the rejection of Claims 24 (and dependent claims thereof) for similar reasons to those given above regarding Claim 1.

Further with respect to Claim 2 (and similarly for Claims 25 and 48), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein a percentage of system resources to be allocated to the process in the class relative to other processes in other classes is based upon the system resource shares associated with the class and the system resource shares associated with the other classes". In rejecting

Claim 2, the Examiner states that Bhatti teaches "specify an number of processes in the server to exclusively support the higher classes and another number of processes in the server" allocated for the lower classes at Bhatti p8 32-47. Applicants have replicated this cited passage below, to show that such passage does not in any way teach any type of higher class exclusivity, as alleged by the Examiner. As stated by Bhatti at Col. 8, lines 32-47:

> "The connection manager 114 is used to determine which access request in which queue is to be accepted next for processing based on the predetermined scheduling policy. The scheduling policy specifies criteria or guidelines for the connection manager 114 to decide which of the queues 111-111n should be serviced next, and thus which request will be passed next to the server application 53. The scheduling policy will be described in more detail below.
>
> In one embodiment, each of the queues 111-111n queues the classified access requests it receives. In another embodiment, each of the queues 111-111n is simply a buffer or storage means that stores the classified access requests it receives.
>
> As described above, the scheduling policy may be selected from a number of known scheduling policies."

As can be seen, the passage cited by the Examiner in rejecting Claim 2 describes a connection manager and scheduling policy co-action for determining which access request in which queue to process next. There is simply no teaching or suggestion in this cited passage of the claimed feature of "wherein a percentage of system resources to be allocated to the process in the class relative to other processes in other classes is based upon the system resource shares associated with the class and the system resource shares associated with the other classes". This claimed feature advantageously determines a *percentage of system resources to be allocated a process.*

In addition, the Examiner's assertion, while not found in the cited passage of Bhatti, fails to establish this claimed feature. The assertion merely specifies a number of

*processes* to support a *class*, but does not teach or otherwise suggest *allocation of system resources to a process*, as claimed, nor the use of system resource shares associated with a class and the other classes as a part of the percentage of process system resource determination. Quite simply, classes and processes are different, and even assuming arguendo that Bhatti teaches allocation of processes to a class, that still does not teach or suggest allocation of system resources to a process, as claimed. Thus, Claim 2 (and similarly for Claims 25 and 48) is further shown to not be obvious in view of the cited references.

Further with respect to Claim 6 (and similarly for Claim 29), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein each class of the plurality of predefined classes has a minimum and a maximum resource limit, the minimum and maximum resource limits defining a minimum and a maximum amount of a resource that may be allocated to the class as a percentage of the resource". This claimed feature advantageously provides a window of resource limits for a class, bounded by a minimum resource limit and a maximum resource limit to thereby define both a minimum and maximum amount of resource that may be allocated to the class. In rejecting Claim 6, the Examiner states that Bhatti teaches "the scheduling policy may specify an number of processes in the server ... to exclusively support the higher classes and another number of processes in the server" assigned as available resources for the lower class at Bhatti p8 32-44 which correspond to the recitation of a minimum and maximum resource limit allocable for each class. Applicants show error in this assertion. Claim 6 is directed to having a minimum and maximum resource limit for a given class. The Examiner's assertion is with respect to two types of classes (higher classes and lower classes), and the specifying of a number of processes to support these two types of classes. This is in contrast to Claim 6, which recites upper and lower bounds for a *given class*. The cited reference does not teach or suggest resource bounds (minimum and maximum) *for a given class*. Thus, Claim 6 (and similarly for Claim 29) is further shown to not be obvious in view of the cited references.

Further with respect to Claim 8 (and similarly for Claim 31), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein performing workload management with respect to other classes within a same tier comprises

determining a percentage goal for the process as a function of a number of system resource shares associated with the class in which the process is classified divided by a total number of shares allocated to active classes in the same tier as the class in which the process is classified". As can be seen, this claim recites determination of a percentage goal for a given process. Two items are used in such determination – (i) a number of system resource shares associated with the class in which the process is classified, and (ii) a total number of shares allocated to active classes in the same tier. For example, if the process class had 5 system resource shares and the total number of shares allocated to active classes in the same tier was 25, the percentage goal for the process would be 5 divided by 25 or 20%.

In rejection Claim 8, the Examiner cites Ferguson p3 39-47 as teaching this claimed percentage goal determination. Applicants show that there, Ferguson states:

> "In the preferred embodiment, the workload manager also priority rates or orders the classes in accordance with the current class performance indices so that a transaction of a class which is performing more poorly (as judged by its *class performance index, which in the preferred embodiment is the ratio of the current average class response time and the class response time goal*) gets a higher dispatch priority at the server to which it is routed than a transaction of a class that is performing better."

As can be seen, this passage mentions a class performance index, which is the ratio of the current average class response time and the class response time goal. It thus teaches using an actual and desired response time *of a given class* to determine the performance index of such class. This is in contrast to using system resource shares for all classes in a tier, as claimed.

In addition, this determination by Ferguson is with respect to response time of a class, and not with respect to system resource shares of a class, as claimed. A response time is merely an indication of the elapsed time taken to respond to a request. It is not a system resource or system resource share.

Thus, Claim 8 (and similarly for Claim 31) is further shown to not be obvious in view of the cited references as there are missing claimed elements not taught or suggested by the cited references.

Further with respect to Claim 9 (and similarly for Claim 32), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein each process in each of the plurality of predefined classes has an associated priority component that is used, along with the system resource shares of the class, to calculate a resource allocation priority". This claimed feature advantageously calculates a resource allocation priority using two items – (i) a priority component of a process, and (ii) the system resource shares of the class. In rejecting Claim 9, the Examiner states that Ferguson teaches this at pp2-3. Applicants have thoroughly reviewed this passage and cannot find any teaching or suggestion of using both a priority component *of a process* and system resource shares *of a class* to calculate a resource allocation priority. If the rejection of Claim 9 is maintained, the Examiner is respectfully requested to point out with specificity where these two items ((i) and (ii) listed above) are discussed. In the absence thereof, Applicants show that the Examiner has failed to properly establish a prima facie showing of obviousness with respect to Claim 9, as the Examiner has merely made a general allegation that the claim language is taught, but has not made any specific showing of such teaching in the cited references.

Further with respect to Claim 10 (and similarly Claim 33), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein performing workload management with respect to other classes within the same tier comprises determining a resource allocation priority range for the class based on a minimum resource usage limit, a maximum resource usage limit, and a resource usage goal". As can be seen, this claim recites a determination of a resource allocation *priority range* for a class, and this class priority range is determined using three items – (i) a minimum resource usage limit, (ii) a maximum resource usage limit, and (iii) a resource usage goal. In rejecting Claim 10, the Examiner merely states "see the discussion of claim 8 supra". In the discussion of Claim 8, the Examiner cites Ferguson p3 39-47. This passage merely teaches using an actual and desired response time *of a given class* (i.e. it only uses two items) to determine the performance index of such class. A *performance* index is not a

*priority* range, as claimed, but rather is an indication of whether response time goal is being met. For example, if the actual response time and response time goal were both 5, the performance index would be 5/5 or 1, indicating that the response time goal was being met. If the actual response time was 10 and the response time goal were 5, the performance index would be 10/5 or 2, indicating that the response time goal was not being met. In contrast, Claim 10 is directed to determining a resource allocation *priority range*. The Ferguson performance index is merely an indication of whether response time goals are being met, and has nothing whatsoever to do with determining any type of priority, much less the claimed resource allocation priority range. Thus, Claim 10 (and similarly Claim 33) is further shown to not be obvious in view of the cited references.

Applicants further traverse the rejection of Claim 11 (and similarly for Claim 34) for similar reasons to those given above regarding Claim 10 (of which Claim 11 depends upon).

Applicants further traverse the rejection of Claim 11 (and similarly for Claim 34) by showing that none of the cited references teach or suggest the claimed feature of "wherein performing workload management with respect to other classes within the same tier further comprises adjusting the resource allocation priority range for the class based on a comparison of an actual amount of the resource being utilized by the class to the maximum resource usage limit, the minimum resource usage limit, and the resource usage goal". As can be seen, Claim 11 recites a step of adjusting the resource allocation priority range for a class. The Examiner has not alleged any teaching or suggestion of this claimed step. Therefore, the Examiner has failed to properly establish a prima facie showing of obviousness with respect to Claim 11. In rejecting claims under 35 U.S.C. Section 103, the examiner bears the initial burden of presenting a prima facie case of obviousness. *In re Oetiker*, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992). Only if that burden is met, does the burden of coming forward with evidence or argument shift to the applicant. Id. "A prima facie case of obviousness is established when the teachings from the prior art itself would appear to have suggested the claimed subject matter to a person of ordinary skill in the art." *In re Bell*, 991 F.2d 781, 782, 26 USPQ2d 1529, 1531 (Fed. Cir. 1993) (quoting *In re Rinehart*, 531 F.2d 1048, 1051, 189 USPQ 143, 147 (CCPA 1976)). If the examiner fails to establish a prima facie case, the

rejection is improper and will be overturned. *In re Fine*, 837 F.2d 1071, 1074, 5 USPQ2d 1596, 1598 (Fed. Cir. 1988). To establish prima facie obviousness of a claimed invention, all of the claim limitations must be taught or suggested by the prior art. MPEP 2143.03. As the Examiner has not alleged a showing of all the claimed steps, a prima facie case of obviousness has not been made and the burden has not shifted to Applicants to rebut an obviousness assertion. Therefore, Claim 11 (and similarly for Claim 34) is further shown to have been erroneously rejected as there are missing claimed elements not taught or suggested by the cited references.

Applicants further traverse the rejection of Claim 12 (and similarly for Claim 35) for similar reasons to those given above regarding Claim 8 (of which Claim 12 depends upon).

Further with respect to Claim 12 (and similarly Claim 35), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein if the percentage goal is below a minimum resource usage limit, the class is favored for additional usage of the system resource, and wherein if the calculated percentage goal is above a maximum resource usage limit, the class is not favored for additional usage of the system resource". As can be seen, Claim 12 recites that if the *percentage goal* is below a minimum limit, the class is favored for additional system resource usage, and if the percentage goal is above a maximum limit, the class is not favored for additional system resource usage. In contrast, the passage cited by the Examiner in rejecting Claim 12 (Ferguson p3 40-47) teaches that if a transaction is performing poorly (i.e. the actual response time is greater than the response time goal), the transaction gets a higher dispatch priority than a transaction that is performing better. Giving a transaction a higher dispatch priority means that it gets dispatched sooner, not that it gets more or less system resource usage, as claimed.

In addition, Claim 12 recites a determination of whether a given value (percentage goal) is *both* below some threshold (minimum limit) *and* above some threshold (maximum limit). The cited passage merely teaches a single determination for a transaction – whether it is performing poorly. Two thresholds, as claimed, are not used in this determination.

Therefore, Claim 12 (and similarly Claim 35) is shown to not be obvious in view of the cited references as there are missing claimed elements not taught or suggested by the cited references.

Applicants further traverse the rejection of Claim 13 (and similarly for Claim 36) for similar reasons to those given above regarding Claim 12 (of which Claim 13 depends upon).

Further with respect to Claim 13 (and similarly for Claim 36), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein the minimum resource usage limit and maximum resource usage limit are retrieved from a share/tier profile storage". In rejecting Claim 13, the Examiner states that Ferguson teaches "System state information" and more specifically, "resource usage statistics" in connection with each class, p6 14-28. Applicants have reviewed this cited passage, and it does not teach or suggest any type of share/tier profile storage. Rather, it discusses details of an Arrival Algorithm, including recording an arrival time, computing a value P(j,l) which is the estimated new value of the class performance index. In contrast, the claimed invention is directed to retrieval of minimum and maximum resource usage *limits*. As described in Applicants' Specification at page 18, line 20 – page 19, line 22, these limits designate the minimum and maximum percentages of system resources that are to be allocated to a class. Even if the cited passage taught resource usage statistics, which is does not, usage statistics have no bearing or relationship to resource usage *limits*, as claimed. Thus, Claim 13 (and similarly for Claim 36) is further shown to have been erroneously rejected as there are missing claimed elements not taught or suggested by the cited references.

Further with respect to Claim 14 (and similarly for Claim 37), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein each of the plurality of predefined classes includes an associated absolute maximum resource usage limit, and wherein if the resource allocation for processes in a class exceed an absolute maximum resource usage limit for the class, additional resource allocation to the class is suspended". As can be seen, additional resource allocation to a class is *suspended* if resource allocation for processes in the class exceeds an absolute maximum resource usage limit for the class. In rejecting Claim 14, the Examiner cites Bhatti p3 40-

55 as teaching "classification system 100 can also reject requests ... to handle potential overload conditions". Applicants show three-fold error in the rejection of Claim 14. First, Claim 14 is directed to *suspension of resource allocation*, not to a rejection of a request. Suspension and rejection are two very different types of actions. Secondly, Claim 14 recites that the classes have yet another characteristic associated with them in addition to the system resource shares and tier. This additional class characteristic is an absolute maximum resource usage limit. The cited reference does not teach or suggest, nor has the Examiner alleged any teaching or suggestion of, a class having an absolute maximum resource usage limit associated with it, as claimed. Thirdly, the cited reference does not teach or suggest any type of determination of whether the *resource allocation for processes in a class* exceed some threshold, as claimed. Rather, the reference merely mentions rejection of a request to handle potential overload conditions or to provide better performance. It also states that the rejection policy can be based on queue length, response time, or any other kind of server load metrics such as CPU utilization or file system activity. Such a general statement does not suggest the particular details as claimed - determining whether the *resource allocation for processes in a class* exceed some threshold, as claimed. Thus, Claim 14 (and similarly for Claim 37) is shown to not be obvious in view of the cited references as there are missing claimed elements not taught or suggested by the cited references.

Further with respect to Claim 15 (and similarly for Claim 38), Applicants show that none of the cited references teach or suggest the claimed step of "determining a class priority adjustment for each of the plurality of predefined classes, wherein the class priority adjustment is determined based on a delta value computed for every predetermined time increment". Nor has the Examiner alleged any such teaching or suggestion. The Examiner merely states in rejecting Claim 15 "see the claim 11 discussion". As to the Claim 11 discussion, the Examiner merely states "see the discussion of claim 8 supra". In the Claim 8 discussion, there is no assertion of any teaching or suggestion of determining a class priority adjustment, where the class priority adjustment is determined based on a delta value computed for every predetermined time increment. Thus, a prima facie case of obviousness has not been made with respect to Claim 15 (and similarly for Claim 38).

Applicants further traverse the rejection of Claim 16 (and similarly for Claim 39) for similar reasons to those given above regarding Claim 15 (of which Claim 16 depends upon).

Further with respect to Claim 16 (and similarly for Claim 39), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein the delta value is a difference between an average resource usage over a specified time interval and a resource usage for a last predetermined time increment". In rejecting Claim 16, the Examiner states that Ferguson teaches at p3 17-20 "scheduling priority for each class of transactions and by dynamically adjusting these scheduling priorities in accordance with the response time dissatisfaction performance index". Applicants show that this dissatisfaction performance index is the ratio of the current average response time and the response time goal (Ferguson Col. 3, lines 41-44), and is not a difference between an average resource usage over a specified time interval and a resource usage from a last predetermined time increment. Even more fundamentally, Ferguson's performance index is determined using an actual and desired value (current average response time and response time *goal*), whereas the claimed delta value is determined using two actual resource usage values (average resource usage over a specified time interval, and resource usage for a last predetermined time increment). Thus, Claim 16 (and similarly for Claim 39) is further shown to not be obvious in view of the cited references as there are missing claimed elements not taught or suggested by the cited references.

Applicants further traverse the rejection of Claim 17 (and similarly for Claim 40) for similar reasons to those given above regarding Claim 10 (of which Claim 17 depends upon).

Further with respect to Claim 17 (and similarly for Claim 40), Applicants show that none of the cited references teach or suggest the claimed step of "adjusting a priority of the class by a class priority adjustment value, the class priority adjustment value being determined based on a comparison of actual resource usage to the minimum resource usage limit, the maximum resource usage limit, and a resource usage goal". As can be seen, this claimed step adjusts a class priority by an adjustment value, and this adjustment value is determined based on a comparison of actual resource usage to three other

parameters – (i) a minimum resource usage limit, (ii) a maximum resource usage limit, and (iii) a resource usage goal. In rejecting Claim 17, the Examiner merely states "see the discussions of claims 6-10 supra. None of the Examiner's discussion involving Claims 6-10 make any mention or allegation of any teaching or suggestion of adjusting a class priority based on three different criteria – (i) a minimum resource usage limit, (ii) a maximum resource usage limit, and (iii) a resource usage goal. At best, the cited reference uses two criteria – an actual response time and a response time goal. A response time is just that, a measure of time. It is not a measure of any type of resource utilization limits or goals, including a minimum resource usage limit, and maximum resource usage limit and a resource usage goal. Thus, the Examiner has failed to establish a prima facie showing of obviousness with respect to Claim 17.

Applicants further traverse the rejection of Claim 18 (and similarly for Claim 41) for similar reasons to those given above regarding Claim 17 (of which Claim 18 depends upon).

Further with respect to Claim 18 (and similarly for Claim 41), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein if the actual resource usage is between the maximum resource usage limit and an absolute maximum resource usage limit, the class priority adjustment is set to disfavor the class from being allocated additional amounts of the resource". As can be seen, this claim recites both a maximum resource usage limit and an absolute maximum resource usage limit. This is described in Applicants' Specification at page 21, line 29 – page 22, line 21, and provides both 'soft' and absolute maximum limits. None of the cited references teach or suggest two types of maximum limits, nor has the Examiner alleged any such teaching or suggestion. Thus, a prima facie showing of obviousness has not been made with respect to Claim 18, and the burden has therefore not shifted to Applicants to rebut an obviousness assertion.

Applicants further traverse the rejection of Claim 19 (and similarly for Claim 42) for similar reasons to those given above regarding Claim 17 (of which Claim 19 depends upon).

Further with respect to Claim 19 (and similarly for Claim 42), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein if the

actual resource usage is between the resource usage goal and the maximum resource usage limit and a difference between an average resource usage over a specified time interval and a resource usage for a last predetermined time increment is less than zero, the class priority adjustment is not changed". Nor has the Examiner alleged any such teaching or suggestion. Thus, a prima facie showing of obviousness has not been made with respect to Claim 19, and the burden has therefore not shifted to Applicants to rebut an obviousness assertion.

Applicants further traverse the rejection of Claim 20 (and similarly for Claim 43) for similar reasons to those given above regarding Claim 17 (of which Claim 20 depends upon).

Further with respect to Claim 20, Applicants show that none of the cited references teach or suggest the claimed feature of "wherein if the actual resource usage is between the minimum resource usage limit and the resource usage goal and a difference between an average resource usage over a specified time interval and a resource usage for a last predetermined time increment is greater than zero, the class priority adjustment is incremented by a function of the difference". Nor has the Examiner alleged any such teaching or suggestion. Thus, a prima facie showing of obviousness has not been made with respect to Claim 20, and the burden has therefore not shifted to Applicants to rebut an obviousness assertion.

Applicants further traverse the rejection of Claim 21 (and similarly for Claim 44) for similar reasons to those given above regarding Claim 17 (of which Claim 21 depends upon).

Further with respect to Claim 21 (and similarly for Claim 44), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein if the actual resource usage is below the resource usage goal and the difference between an average resource usage over a specified time interval and a resource usage for a last predetermined time increment is greater than zero, the class priority adjustment is not changed". Nor has the Examiner alleged any such teaching or suggestion. Thus, a prima facie showing of obviousness has not been made with respect to Claim 21, and the burden has therefore not shifted to Applicants to rebut an obviousness assertion.

Applicants further traverse the rejection of Claim 22 (and similarly for Claim 45) for similar reasons to those given above regarding Claim 17 (of which Claim 22 depends upon).

Further with respect to Claim 22 (and similarly for Claim 45), Applicants show that none of the cited references teach or suggest the claimed feature of "wherein if the actual resource usage is below the resource usage goal and the difference between an average resource usage over a specified time interval and a resource usage for a last predetermined time increment is less than or equal to zero, the class priority adjustment is decremented by a function of the difference". Nor has the Examiner alleged any such teaching or suggestion. Thus, a prima facie showing of obviousness has not been made with respect to Claim 22, and the burden has therefore not shifted to Applicants to rebut an obviousness assertion.

Applicants further traverse the rejection of Claim 23 (and similarly for Claim 46) for similar reasons to those given above regarding Claim 7 (of which Claim 23 depends upon).

Further with respect to Claim 23, Applicants show that none of the cited references teach or suggest the claimed feature of "wherein processes in classes in lower priority tiers are only able to obtain access to a resource if processes in classes in higher priority tiers do not use all of the resource or if the processes in each of the classes in the higher priority tier have reached an absolute maximum resource usage limit". Nor has the Examiner alleged any such teaching or suggestion. Thus, a prima facie showing of obviousness has not been made with respect to Claim 23, and the burden has therefore not shifted to Applicants to rebut an obviousness assertion.

With respect to Claim 47 (and dependent claims thereof), Applicants show that none of the cited references teach or suggest classifying a process into a class having two-fold characteristics (system resource shares and tier). As previously discussed, the only mention by Bhatti of process classification is a single-fold priority classification. In contrast, the claimed invention is directed to classifying a process into a class, where the class has *two-fold* characteristics (system resource shares and tier). As every element of Claim 47 is not taught or suggested by the cited reference, it is shown that Claim 47 (and dependent claims thereof) is not obvious in view of the cited references.

Therefore, the rejection of Claims 1-50 under 35 U.S.C. § 103 has been overcome.

## III. Conclusion

It is respectfully urged that the subject application is patentable over the cited references and is now in condition for allowance. The Examiner is invited to call the undersigned at the below-listed telephone number if in the opinion of the examiner such a telephone conference would expedite or aid the prosecution and examination of this application.

DATE: 3/17/04

Respectfully submitted,

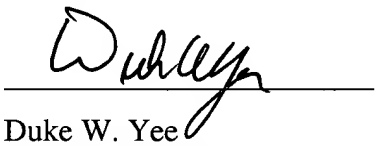

Duke W. Yee
Reg. No. 34,285

Wayne P. Bailey
Reg. No. 34,289
Carstens, Yee & Cahoon, LLP
P.O. Box 802334
Dallas, TX 75380
(972) 367-2001
Attorneys for Applicants